AD-A144 414 RESEARCH ON PARALLELISM IN PROBLEM-SOLVING SYSTEMS(U) 1/1
SRI INTERNATIONAL MENLO PARK CA D E WILKINS 03 JUL 84
AFOSR-TR-84-0595 F49620-79-C-0188

UNCLASSIFIED F/G 12/2 NL

END

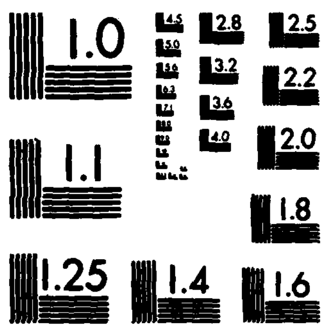

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

SRI International

AFOSR-TR- 84-0595

(5)

ANNUAL SCIENTIFIC REPORT

Period Covering: September 1, 1982 to August 31, 1983

July 3, 1984

By: David E. Wilkins, Computer Scientist

Artificial Intelligence Center
Computer Science and Technology Division

Prepared for:

Air Force Office of Scientific Research
Building 410
Bolling Air Force Base
Washington, D.C. 20332

ATTENTION: Dr. Buchal

AFOSR Contract No. F49620-79-C-0188

SRI Project 8871

SRI International

333 Ravenswood Ave. • Menlo Park, CA [illegible]
415 326-6200 • TWX 910 373-2046 • Telex 334-[illegible]

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S) **AFOSR-TR- 84-0595**
6a. NAME OF PERFORMING ORGANIZATION SRI, International 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research
6c. ADDRESS (City, State, and ZIP Code) Artificial Intelligence Ctr, Computer Science & Technology Div, 333 Ravenswood Ave, Menlo Park CA 94025	7b. ADDRESS (City, State, and ZIP Code) Directorate of Mathematical & Information Sciences, AFOSR, Bolling AFB DC 20332
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR 8b. OFFICE SYMBOL (If applicable) NM	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F49620-79-C-0188
8c. ADDRESS (City, State, and ZIP Code) Bolling AFB DC 20332	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. 61102F; PROJECT NO. 2304; TASK NO. A7; WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification)
'RESEARCH ON PARALLELISM IN PROBLEM-SOLVING SYSTEMS',

12. PERSONAL AUTHOR(S)
David E. Wilkins

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Annual	FROM 1/9/82 TO 31/8/83	3 JUL 84	6

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This report describes the progress to date on work to date under the contract. Research on planning and problem-solving systems was begun at SRI International in September 1979. Progress has been described in detail in three annual reports (1980, 1981, and 1982). The main task of this research program is to develop powerful methods of representing, generating and executing hierarchical plans that contain parallel actions. Execution involves monitoring the state of the world and possibly replanning if things do not proceed as expected. Two different approaches to these problems are being pursued under this contract. The first is heuristic; it involves building an actual computer program that provides a representation from which it then generates plans. This approach comprises the majority of the effort on this project. The second approach is to investigate the theoretical foundations of planning. This will not, in all likelihood, result in a program, but it will formalize the planning problem and one solution to it.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Robert N. Buchal	22b. TELEPHONE (Include Area Code) (202) 767- 4939	22c. OFFICE SYMBOL NM

DD FORM 1473, 84 MAR — 83 APR edition may be used until exhausted. All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE
UNCLASSIFIED

84 08 07 137

INTRODUCTION

This report describes the progress to date on work described in SRI's Proposal ECU 82-060, "Research on Parallelism in Problem-Solving Systems". This two year proposal was negotiated as Modification No. 9 of AFOSR Contract F49620-79-C-0188. This modification provided funding for the Year One portion of the work; Year Two was treated as an option that could be exercised at the contracting officer's discretion by September 1, 1983. This letter summarizes what has been accomplished during the first year.

Research on planning and problem-solving systems was begun at SRI International in September 1979 under AFOSR sponsorship (SRI Project 8871; Contract No. F49620-79-C-0188). Progress has been described in detail in three annual reports (1980, 1981, and 1982). The main task of this research program is to develop powerful methods of representing, generating, and executing hierarchical plans that contain parallel actions. Execution involves monitoring the state of the world and possibly replanning if things do not proceed as expected. Two different approaches to these problems are being pursued under this contract. The first is heuristic; it involves building an actual computer program that provides a representation from which it then generates plans. This approach comprises the majority of the effort on this project. The second approach is to investigate the theoretical foundations of planning. This will not, in all likelihood, result in a program, but it will formalize the planning problem and one solution to it.

I SIPE

A principal goal of our research in planning and plan execution is the development of a heuristic system that can plan an activity and then monitor the execution of that plan. Over the last two years we have designed and implemented such a system, SIPE, (System for Interactive Planning and Execution Monitoring). The basic approach to planning is to work within the hierarchical-planning paradigm, representing plans in procedural networks -- as has been done in NOAH [2] and other systems. Several extensions of previous planning systems have been implemented, including the development of a perspicuous formalism for describing operators and objects, the use of constraints for the partial description of objects, the creation of mechanisms that permit concurrent exploration of alternative plans, the incorporation of heuristics for reasoning about resources, and the creation of mechanisms that make it possible to perform deductions.

This year we have begun using the planning of assembly tasks for a robotics workstation as a motivating domain. This choice is based on those characteristics that are suited to automatic planning. It is easy to define hierarchical levels and metalevels in the problem, which is reasonable to represent in the formalisms we have developed and is difficult enough to provide a challenge. Plans using multiple arms can be generated, thus further challenging the parallel-planning capabilities of our system. The idea is not to produce a working system that will interact with the robotics workstation here, but rather to test our ideas on the real problems provided by the workstation.

This domain also furnishes us with many real-world considerations. We would like to investigate planning techniques that run in approximately real time. There is an element of uncertainty in the robotics domain (e.g., the arm does not always put down an object in exactly the intended location) that will stress execution-monitoring capabilities and require operations for checking the state of the world. These checking operations may be expensive (e.g., taking pictures with the vision module and analyzing them), so they must be carefully planned for by the system. These checks will also require conditionals in the plan.

We intend to model robotics assembly tasks adequately. Extending our present system, SIPE, to operate in this domain requires finding solutions to many representation problems. Our formalism will have to be expanded to deal with sets, uncertainty, conditionals, loops, and planning to acquire information. The first part of the year was spent investigating other formalisms and approaches to see whether there might be a more suitable one for this task.. Logical formalisms seem advantageous for metaplanning [3]. Both PROLOG and other logical formalisms (such as the one Rosenschein is developing as part of this project) have been investigated to determine whether they are appropriate and offer advantages for a planning system. This research led to the conclusion that most approaches based on logic still suffer from inefficiency because of an inability to control the possible deductions. A number of researchers are exploring such an approach (here at SRI, at Stanford, and at other centers), while the heuristic approach used in SIPE is unique and promising. Therefore, the decision was made to continue the approach used in SIPE, and to attempt an expansion of the program to cover more aspects of planning.

The decision to continue developing SIPE means accepting the assumptions behind SIPE that imply a model approach to representation [4]. We have explored the notion of doing metaplanning within SIPE, but it appears that what we need to say about plans at a metalevel cannot fit easily into the model approach. For example, it would be difficult to represent explicitly every property of a plan, a failed search branch, an operator, or a constraint that we might want to reason about at the metalevel. As a result, we view the other areas mentioned below as more promising areas upon which to concentrate our effort during the remainder of the project. It is possible that some metalevel reasoning will be included in the improved execution-monitoring and resource-reasoning capabilities described below.

Later in the year we designed additions to SIPE that will accommodate many of the features needed in the robotics domain. In particular, representations have been designed for conditional plans and information-gathering operators. The latter means the incorporation of uncertainty into the system, which is a fundamental change. We do not intend to investigate geometric modeling or spatial reasoning, so the new additions have been designed to include "hooks" for incorporating special-purpose subsystems for geometric modeling or spatial reasoning. None of these new designs was implemented during the first year; we expect to accomplish that during the second year of the project. The problems in having uncertainty in the planner are still not all resolved, so the design process is still in progress.

The current resource mechanism in SIPE is a major advance over previous planning systems, yet is only a simple scheme for specifying what happens *during* the performance of an action. It is limited in that it does not allow for interesting uses of shared resources and makes no provision for more general statements about what happens during an action. We have studied this problem to determine what general statements one would like to make about what happens during the performance of actions. Such statements can then be used to prevent harmful interactions in parallel branches of the plan. This will require extensions of the interaction analysis techniques currently used by SIPE. Fortunately, the type of reasoning required can make direct use of the deductive machinery and constraint satisfaction algorithms already available in SIPE. These extensions have also not been implemented in SIPE.

We are working on the problem of developing better execution-monitoring techniques and a greater capability to replan when things do not go as expected. We have accomplished the initial design of some general techniques as well as the ability to incorporate error-handling instructions within operators. The central ideas, outlined below, have not yet been implemented. We have finished the design and expect to do the implementation during the second year of the project.

The intent is for SIPE to handle two different kinds of inputs during execution: (1) Predicates may be specified at any point during execution. These may be anything, no matter how unexpected, in the domain language being used by the planner. They may come from the user or from a high-level sensor such as another computer. (2) Information may be received automatically from sensors. In general, one can predict what might be received (e.g., the sensors in the fingers of the arm may tell us the finger separation is zero, which means that the hand is empty -- but they won't tell us the table has collapsed).

In both cases, there may not be a full specification as to what has changed, so we will need to deduce some predicates. This involves expanding SIPE's deductive capability. There will be a general replanning capability that handles Case 1, but it will be based more on brute-force than cleverness. To handle Case 2, we have designed an extension of the Operator Description Language to represent instructions within operators regarding the actions to be taken after certain foreseeable errors have occurred. These instructions should be encoded in the operator, since the proper context exists there for encoding this knowledge. If this is done separately, the context for use of the instructions will somehow have to be described. Knowledge embodied in these instructions can prevent us from applying general abilities, since they tell us that the scope of possible effects is limited. In this way, proper corrective actions can be generated quickly.

II THEORETICAL FOUNDATIONS

Most work on robot planning and problem-solving is done against the background of an implicit but unarticulated theory of rational action. Roughly stated, this theory assumes a rational agent who attempts to maintain a state of consistency between his intentions (plans) and his beliefs and goals; the agent will perform (and intend to perform) those actions he believes will achieve his goals. We have developed a model of the cognitive agent in which the principle of rationality is formalized and studied abstractly. The model is quite general; beliefs may be certain or uncertain, goals are treated as a special case of preferences, and multiple agents can be handled in a natural manner.

The model has immediate value as an analytical tool and a way of integrating such topics as belief revision, execution monitoring, replanning, and various types of goals (e.g., maintenance and prevention) in a common theoretical framework. Furthermore, it raises interesting algorithmic questions about how a robot might compute rational behavior in real time. As was the case with our logical formalisms for reasoning about plans, the formal model of rationality, as it stands now, is general-purpose and does not offer much support in handling important special cases.

III PUBLICATIONS

Part of the project time during the first year was spent writing and revising a major paper describing work done on the project. This paper, entitled "Domain-independent Planning: Representation and Plan Generation", appeared in the Artificial Intelligence Journal in the April 1984 issue. [4]

A condensed version of this paper appeared as a long paper in the 1983 Proceedings of the International Joint Conference on Artificial Intelligence. This project also supported the presentation of this paper by David Wilkins at the conference in Karlsruhe, Germany in August 1983.

IV SUMMARY

After investigating new formalisms, we have decided to expand the current system, since it constitutes a unique and interesting approach to the problem. We have discovered and clarified complications that arise when SIPE is used for metaplanning, and our future efforts will concentrate on other areas. A major paper has been written and accepted by the Artificial Intelligence Journal. We have designed representations for conditional plans, information-gathering operators, and error-handling instructions that we intend to incorporate into the system during the last year of the project. With regard to the theoretical foundations of planning, we have developed a model of the cognitive agent in which the principle of rationality is formalized.

REFERENCES

1. Rosenschein, S., "Plan Synthesis: A Logical Perspective", *Proceedings IJCAI--81*, Vancouver, British Columbia, 1981, pp. 331--337.

2. Sacerdoti, E., *A Structure for Plans and Behavior*, Elsevier, North-Holland, New York, New York, 1977.

3. Weyhrauch, R., "Prolegomena To a Theory of Mechanized Formal Reasoning", in *Readings in Artificial Intelligence*, Nilsson and Webber, eds., Tioga Publishing, Palo Alto, California, 1981, pp. 173--191.

4. Wilkins, D., "Domain-independent Planning: Representation and Plan Generation", *Artificial Intelligence 22*, April 1984, pp. 269-301.

END

FILMED

9-84

DTIC